Temperature scaling law for quantum annealing optimizers

Tameem Albash,[1,2] Victor Martin-Mayor,[3,4] and Itay Hen[1,2]

[1] *Information Sciences Institute, University of Southern California, Marina del Rey, California 90292, USA*
[2] *Department of Physics and Astronomy and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, California 90089, USA*
[3] *Departamento de Física Teórica I, Universidad Complutense, 28040 Madrid, Spain*
[4] *Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Zaragoza, Spain*

Physical implementations of quantum annealing unavoidably operate at finite temperatures. We argue that a fixed finite temperature prevents annealers from functioning as competitive scalable optimizers and show that to serve as optimizers annealer temperatures must be appropriately scaled down with problem size. We derive a temperature scaling law dictating that temperature must drop at the very least in a logarithmic manner but also possibly as a power law with problem size. We corroborate our results by experiment and simulations and discuss the implications of these to practical annealers.

arXiv:1703.03871v1 [quant-ph] 10 Mar 2017

Introduction.— Despite the excitement brought on by recent technological breakthroughs that have made programmable quantum annealing (QA) [1–5] optimizers consisting of thousands of quantum bits commercially available, quantum annealing devices have so far failed to deliver on their promise to serve as useful optimizers, i.e., to find bit assignments that minimize the energy, or cost, of discrete combinatorial optimization problems faster than possible classically. Thus far, no examples (neither experimental nor theoretical) of practical relevance have been found to indicate a superiority of QA optimization over traditional methods [6–13]. While recent years have witnessed a shift in the role assigned to quantum annealers from optimizers to samplers [14–16], there is still a major ongoing effort to build larger, more densely connected QA devices, in the hope that the capability to embed larger optimization problems would eventually reveal the coveted quantum speedup [17–21].

In this work, we show that for QA devices to serve as optimizers, their temperature must be scaled down with increasing problem size. We analyze the theoretical performance of *ideal fixed-temperature* quantum annealers for optimization. We show that even in the case where annealers are assumed to *thermalize instantly* (rather than only in the infinite runtime limit), the energies, or costs, of their output configurations would be computationally trivial to achieve (in a sense that we explain). We further derive a scaling law for quantum annealing optimizers and provide corroboration of our analytical findings by experimental results obtained from the commercial D-Wave 2X quantum annealing processor [22–26] as well as numerical simulations. We discuss the implications of our results for both the engineering requirements of QA devices and their benchmarking.

Fixed-temperature quantum annealers.— In the adiabatic limit, closed-system quantum annealers are guaranteed to find a ground state of the cost function, or final Hamiltonian, they are to solve. The adiabatic theorem of quantum mechanics ensures that the overlap of the final state of the system with the ground state manifold of the target cost function, whose Hamiltonian we denote by H, approaches unity as the duration of the process increases [27, 28]. For physical quantum annealers that operate at finite temperatures ($T > 0$), there is no equivalent guarantee of reaching the ground state with high probability. In the limit of very long runtimes, an ideal finite-temperature quantum annealer is expected to sample the Boltzmann distribution of the final Hamiltonian at the annealer temperature [29].

In what follows, we argue that even instantly-thermalizing quantum annealers are severely limited as optimizers due to their finite temperature. For concreteness, we will consider in our analysis annealers for which i) the number of couplers scales linearly with the number of qubits N [30], ii) the coupling strengths are discretized and are bounded independently of problem size, and iii) the scaling of the free energy with problem size is not pathological, i.e., that our system is not tuned to a critical point. Other than the above rather standard assumptions, our treatment will be general (we discuss the performance of quantum annealers when some of these conditions are lifted later on). For clarity, it is useful to have in mind optimization problems written in terms of a Hamiltonian of the Ising-type

$$H = \sum_{\langle ij \rangle} J_{ij} s_i s_j + \sum_i h_i s_i \,, \tag{1}$$

where $\{s_i = \pm 1\}$ are binary Ising spin variables that are to be optimized over, $\{J_{ij}, h_i\}$ are the coupling strengths between connected spins and external biases, respectively, and $\langle ij \rangle$ denotes the underlying connectivity graph of the model. The discussion that follows however is not restricted to any particular model.

Under the above assumptions, the ground state energies, denoted E_0, of any given problem class, scale linearly with increasing problem size (i.e., the energy is an extensive property as is generically expected from physical systems) while the minimal gap $\Delta = E_1 - E_0$ remains fixed. From our assumptions it follows then [31] that the

thermal expectation values of the intensive energy

$$\langle e\rangle_\beta = \langle H\rangle_\beta / N\,, \tag{2}$$

and specific heat

$$c_\beta = \partial\langle e\rangle_\beta/\partial\beta = -N\left[\langle e^2\rangle_\beta - \langle e\rangle_\beta^2\right]\,, \tag{3}$$

remain finite as $N \to \infty$ for any fixed inverse-temperature $\beta = 1/T$. The intensive energy is discretized in steps of Δ/N, yet its statistical dispersion $\sigma_\beta(e) = \sqrt{-c_\beta/N}$ is much larger. Treating e as a stochastic variable, for large enough values of N it can be treated as a continuous variable as the ratio of discretization versus dispersion is $\sqrt{-c_\beta\Delta^2/N}$ decaying to zero for large N. From the Boltzmann distribution it follows that the probability density of e goes as $p_\beta(e) = {Z_\beta}^{-1}\mathrm{e}^{N(s(e)-\beta e)}$, where $Z_\beta = \sum_n g_n \mathrm{e}^{-\beta E_n}$ is the partition function, g_n is the degeneracy of the n-th level, i.e., the number of microstates with $H(\{s_i\}) = E_n$, satisfying $2^n = \sum_{n\geq 0} g_n$, and $s(e)$ is the entropy density [32]. The linear combination $\Psi_\beta(e) = s(e) - \beta e$ plays the role of a large-deviations functional for e. The most probable value of e, which we denote by e^*, is given by the maximum of Ψ_β. Solving $\Psi'_\beta(e^*) = 0$, we find [33]

$$\beta = \left.\frac{\partial s}{\partial e}\right|_{e=e^*}\,. \tag{4}$$

Close to e^*, Ψ_β can be Taylor-expanded as $\Psi_\beta(e) \approx \Psi_\beta(e^*) - \frac{|\Psi''_\beta(e^*)|}{2}(e-e^*)^2$, from which it follows that

$$p_\beta(e) \approx \frac{\mathrm{e}^{\Psi_\beta(e^*)}}{Z_\beta}\exp\left[-\frac{N|\Psi''_\beta(e^*)|}{2}(e-e^*)^2\right]\,. \tag{5}$$

The probability density is thus approximately Gaussian in the vicinity of e^*. Moreover, in the limit of large N, we find

$$\langle e\rangle_\beta = e^* \quad \text{and} \quad c_\beta = \frac{-1}{|\Psi''_\beta(e^*)|}\,. \tag{6}$$

Therefore, the probability of finding by Boltzmann-sampling any energy $e < e^*$ (equivalently, $E < e^*N$) is exponentially suppressed in N, scaling in fact as $\exp[-N(\,\Psi_\beta(e^*) - \Psi_\beta(e)\,)]$. We thus arrive at the conclusion that even ideal fixed temperature quantum annealers that thermalize instantaneously to the Gibbs state of the classical Hamiltonian are *exponentially unlikely* to find the ground state since $e^* > e_0 \equiv E_0/N$.

We now corroborate the above derivation by runs on the commercial DW2X quantum annealer [22–25]. To do so, we first generate random instances of differently sized sub-graphs of the DW2X Chimera connectivity graph [34, 35] and run them multiple times on the annealer, recording the obtained energies [36]. Figure 1 depicts typical resultant residual energy $(E - E_0)$ distributions. As is evident, increasing the problem size N 'pushes' the energy distribution farther and farther away from the ground state value, as well as broadening the distribution and making it more gaussian-like. In the inset, we measure the departure of $\langle H\rangle_\beta$ from E_0 and the spread of the energies $\sigma_\beta(H)$ over 100 'planted-solution' [11] instances per sub-graph size as a function of problem size N [37]. For sufficiently large problem sizes, we find that the scaling of $\langle H - E_0\rangle_\beta$ is linear while $\sigma_\beta(H)$ scales as $\sqrt{N}$, consistent with the analytical prediction (the solid lines are linear and square-root fits, respectively, to the largest four sizes). We note though that despite their consistent behavior with our preliminary assumptions it is yet unclear whether the configurations obtained from the DW2X processor are equilibrated.

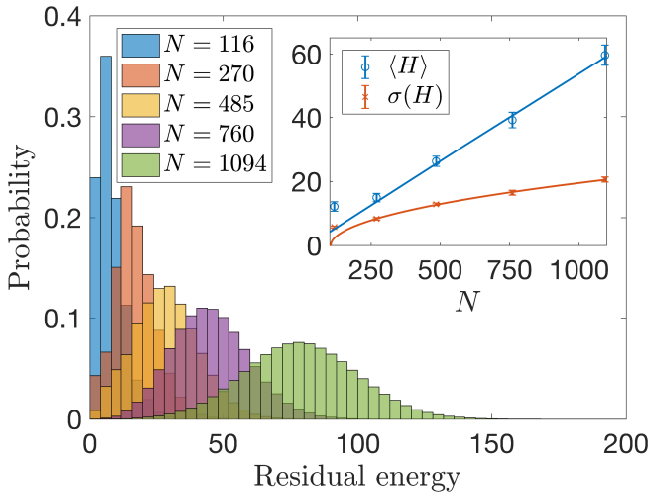

FIG. 1. **Distributions of residual energy, $E - E_0$, from DW2X simulations.** As problem sizes grow, the distributions become more Gaussian-like. **Inset:** Gaussians' mean (blue) and standard deviation (red) as a function of problem size, averaged over 100 instances per size. The solid lines are fits to a straight line and a square root, respectively, taking into accounts all sizes but the smallest.

Given the scaling of the mean and standard deviation, we conclude that fixed-temperature quantum annealers will generate energies e with a fixed distance from e_0, or in terms of extensive energies, configurations obtained from fixed-temperature annealers will have energies concentrated around $E = (1-\epsilon)E_0$ for some $\epsilon > 0$.

One could now ask what the difficulty is for *classical algorithms* to generate energy values in the above range. This question has been recently answered, by the discovery of a polynomial time approximation scheme (PTAS) for spin-glasses defined on a Chimera graph [38] (and which can be easily generalized to any locally connected model), where it has been demonstrated that reaching such energies can be done efficiently [39]. While it is important to note that the scaling of the PTAS with ϵ is not favorable, scaling as $c^{1/\epsilon}$ for some constant c, in practice there exist algorithms (e.g., parallel tempering

that we discuss later on) that are known to scale more favorably than PTAS.

Scaling law for quantum annealing temperatures.— In light of the above, it may seem that quantum annealers are doomed to fail as optimizers as problem sizes increase. We now argue that success may be regained if the temperature of the QA device is appropriately scaled with problem size. Specifically, we address the question of how the inverse-temperature β should scale with N such that there is a probability of at least q of finding the ground state.

An estimate for the required scaling can be given as follows. From the above analysis, it should be clear that the probability of finding a ground state at inverse temperature β will not decay exponentially with system size only if the ground state falls within the variation of the mean energy, specifically if

$$\sigma_\beta(H) = N\sigma_\beta(e) = \sqrt{-Nc_\beta}\,, \tag{7}$$

is comparable to

$$\langle H\rangle_\beta - E_0 = -N\int_\beta^\infty \mathrm{d}\,\beta\; c_\beta\,. \tag{8}$$

The third law of thermodynamics dictates that the specific heat $c_T \equiv \mathrm{d}\langle e\rangle/\mathrm{d}T$ goes to zero when $T \to 0$. Assuming a scaling of the form $c_T \sim T^\alpha$, or equivalently, $-c_\beta \sim \beta^{-\alpha-2}$, gives

$$\sigma_\beta(H) \sim \sqrt{\frac{N}{\beta^{\alpha+2}}} \quad \text{and} \quad \langle H\rangle_\beta - E_0 = \frac{N}{\beta^{\alpha+1}}\,. \tag{9}$$

For a power-law specific heat, it thus follows that the sought scaling is $\beta \sim N^{1/\alpha}$. If on the other hand c_β vanishes exponentially in β, the inverse-temperature scaling will be milder, of the form $\beta \sim \log N$.

To illustrate the above, we next present an analysis of simulations of randomly generated instances on Chimera lattices (we study several problem classes and architectures, see SI). To study the energy distribution generated by a thermal sampler on these instances, we use parallel tempering (PT) [40, 41], a Monte Carlo method whereby multiple copies of the system at different temperatures are simulated [42]. We equilibrate the instances at different temperatures and different problem sizes. In Fig. 2, we show an example of how the energy distribution of a planted-solution instance changes with β. The qualitative behavior is similar to what we observe with increasing problem size, whereby decreasing β (increasing the temperature) pushes the energy distribution to larger energies and makes it more gaussian-like.

The behavior of the specific heat c_β as the inverse-temperature β becomes large is shown in Fig. 3. At large sizes, the scaling becomes $c_\beta \propto \exp(-\Delta\beta)$ as expected (here, $\Delta = 4$ is the gap). Based on our predictions above,

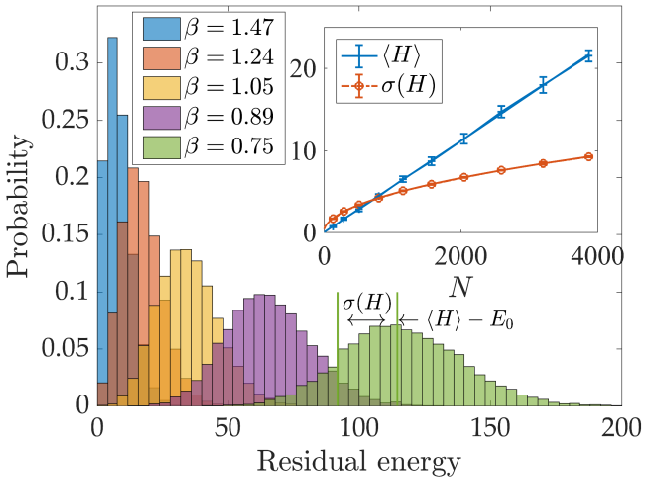

FIG. 2. **Distributions of residual energy, $E - E_0$, from PT simulations.** For a planted-solution instance defined on an $L = 12$ Chimera graph, the distributions become more Gaussian-like as β decreases. For the case of $\beta = 0.75$, the mean residual energy and standard deviation are indicated. **Inset:** Scaling with problem size of the median mean energy and median standard deviation of the energy for $\beta = 1.47$ over 100 instances.

this should mean that if for a fixed q, the minimum β^* such that $p_{\beta^*}(E_0) \geq q$ falls in this exponential regime, then we should observe a scaling $\beta^* \propto \log N$. Indeed, the inset of Fig. 3, which shows simulation results of β^* versus N, exhibits the expected $\log N$ behavior [43].

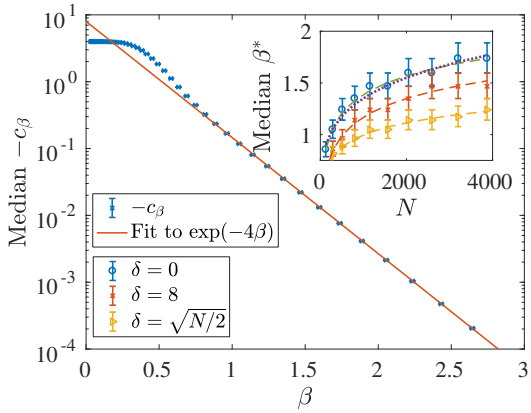

FIG. 3. **Typical specific heat with inverse-temperature.** Behavior of the median specific heat (over 100 instances) for planted-solution instances with inverse-temperature β for $N = 3872$. The behavior transitions from a polynomial scaling with β to an exponential scaling. **Inset:** Typical minimum inverse-temperature required for instances of size N such that the probability of the target energy $E_T = E_0 + \delta(N)$ is at least $q = 10^{-1}$. Also shown are fits to $\log N$ for all three cases and a power-law fit to cN^α that finds $\alpha = 0.19 \pm 0.05$ for the $\delta = 0$ case, which is almost indistinguishable from the logarithmic fit.

While for problem classes with a fixed minimum gap

Δ, one may naively expect c_β to vanish exponentially in general, implying that a logarithmic scaling of β will generally be sufficient as our simulations indeed indicate, it is important to note that two-dimensional spin glasses are known to exhibit a crossover between an exponential behavior to a power law [44–47]. This crossover is characterized by a constant $\theta \approx 1/2$, whereby the discreteness of the gap Δ is evident only for sizes $N^{\theta/2} \ll \beta$. Beyond $N^{\theta/2} \sim \beta$, the $2d$ system behaves as if the coupling distribution is continuous [45, 46] at which point the system can be treated as if with continuous couplings, for which the specific heat c_T scales as T^α with $\alpha_c = 2\nu$ [44], where $\nu = 3.53(7)$ [47]. Therefore, for an ideal quantum annealer operating beyond the crossover, a scaling of $\beta \sim N^{1/(2\nu)\approx 0.14}$ is required. We may thus expect the same crossover to appear for instances defined on the Chimera lattice, which is $2d$-like. Interestingly, for the temperature scaling shown in the inset of Fig. 3, a power-law fit $\beta \sim N^\alpha$ with $\alpha = 0.19 \pm 0.05$ is almost indistinguishable from the logarithmic one, with a power that is consistent with the $2d$ prediction.

Suboptimal metrics for optimization problems.— For many classically intractable optimization problems, when formulated as Ising models, it is crucial that solvers find a true minimizing bit assignment rather than low lying excited states. This is especially true for NP-complete/hard problems [48] where sub-optimal costs generally correspond to violated constraints that *must* be satisfied (otherwise the resultant configuration is nonsensical despite its low energy). Nonetheless, it is plausible to assume the existence of problems for which slightly sub-optimal configurations would still be of value [49]. We thus also study the necessary temperature scaling for cases where the target energies obey $E_{\rm T} \leq E_0 + \delta(N)$ with $\delta(N)$ scaling sub-linearly with problem size. In the inset of Fig. 3, we plot the required scaling of β for $\delta(N) = \text{const}$ and $\delta(N) \propto \sqrt{N}$. In both cases we find that a logarithmic scaling is still essential, albeit with smaller prefactors.

Conclusions and discussion.— We analyzed the behavior of the energy distributions generated by fixed- and variable-temperature ideal quantum annealers with problem size. We have shown that fixed temperature quantum annealers can only sample 'easily reachable' energies in the large problem size limit. We further derived a temperature scaling law to ensure that quantum annealing optimizers find nontrivial energy values with sub-exponential probabilities. We argued that the scaling of the specific heat with temperature controls this scaling: if β lies in the regime where the specific heat scales exponentially with β, then the inverse-temperature of the annealer must scale as $\log N$. However, further considerations need to be taken into account because of possible crossover behavior in the specific heat with temperature and problem size. It would therefore be of interest to study the effect of different architectures that keep the required temperature scaling mild (logarithmic) for large problem sizes as well. The sizes at which a crossover from logarithmic to power-law takes place is highly dependent on the connectivity graph of the QA device. For Chimera graphs, because of their essentially two-dimensional structure, there is a concern that it would lead to a crossover to power law scaling even though simulations suggest a mild logarithmic behavior. Little is known about this crossover in three dimensions or for different architectures, so this concern may not be mitigated by a more complex connectivity graph.

Our results shed important light on previous benchmarking studies that have found no quantum speedups [10, 11, 50–52], identifying temperature as a likely culprit for their unfavorable performance. Furthermore, our analysis is particularly relevant for both the utility as well as the design of future QA devices that have been argued to sample from thermal or close-to-thermal distributions [14], calling their role as optimization devices into question.

Alternative approaches to scaling down the temperature with problem size is the (theoretically) equivalent scaling up of the overall energy scale of the Hamiltonian. However, the rescaling of the total Hamiltonian is also known to be tremendously challenging (for instance, for inductively coupled flux qubits, this would require increasing the currents in the superconducting ring while keeping all other things constant) and may not represent a convenient approach for a scalable architecture. An intriguing alternative approach is to develop quantum error correction techniques to effectively increase the energy scale of the Hamiltonian by coupling multiple qubits to form a single logical qubit [53–58]. Alternatively, quantum error correction techniques can also be used to effectively decouple the system from the bath [59–62].

While our arguments above indicate that fixed-temperature quantum annealers may not be suitable as optimizers, it is important to note that the current study does not pertain to the usage of quantum annealers as *samplers*, where the objective is to sample from the Boltzmann distribution. The latter objective is known to be very difficult task (it is #P-hard [63–65]) and little is known about when or if quantum annealers can provide an advantage in this regard [66].

Acknowledgements.— The computing resources were provided by the USC Center for High Performance Computing and Communications. TA was supported under ARO MURI Grant No. W911NF-11-1-0268, ARO MURI Grant No. W911NF-15-1-0582, and NSF Grant No. INSPIRE- 1551064. V.M.-M. was partly supported by MINECO (Spain), through Grant No. FIS2015-65078-C2-1-P.

[1] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, "Quantum annealing: A new method for minimizing multidimensional functions," Chemical Physics Letters **219**, 343–348 (1994).

[2] J. Brooke, D. Bitko, T. F., Rosenbaum, and G. Aeppli, "Quantum annealing of a disordered magnet," Science **284**, 779–781 (1999).

[3] Tadashi Kadowaki and Hidetoshi Nishimori, "Quantum annealing in the transverse Ising model," Phys. Rev. E **58**, 5355 (1998).

[4] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser, "Quantum Computation by Adiabatic Evolution," arXiv:quant-ph/0001106 (2000).

[5] Giuseppe E. Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car, "Theory of quantum annealing of an Ising spin glass," Science **295**, 2427–2430 (2002).

[6] A. P. Young, S. Knysh, and V. N. Smelyanskiy, "Size dependence of the minimum excitation gap in the quantum adiabatic algorithm," Phys. Rev. Lett. **101**, 170503 (2008).

[7] A. P. Young, S. Knysh, and V. N. Smelyanskiy, "First-order phase transition in the quantum adiabatic algorithm," Phys. Rev. Lett. **104**, 020502 (2010).

[8] Itay Hen and A. P. Young, "Exponential complexity of the quantum adiabatic algorithm for certain satisfiability problems," Phys. Rev. E **84**, 061152 (2011).

[9] E. Farhi, D. Gosset, I. Hen, A. W. Sandvik, P. Shor, A. P. Young, and F. Zamponi, "Performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs," Phys. Rev. A **86**, 052334 (2012), (arXiv:1208.3757).

[10] Troels F. Rønnow, Zhihui Wang, Joshua Job, Sergio Boixo, Sergei V. Isakov, David Wecker, John M. Martinis, Daniel A. Lidar, and Matthias Troyer, "Defining and detecting quantum speedup," Science **345**, 420–424 (2014).

[11] Itay Hen, Joshua Job, Tameem Albash, Troels F. Rønnow, Matthias Troyer, and Daniel A. Lidar, "Probing for quantum speedup in spin-glass problems with planted solutions," Phys. Rev. A **92**, 042325– (2015).

[12] Sergio Boixo, Tameem Albash, Federico M. Spedalieri, Nicholas Chancellor, and Daniel A. Lidar, "Experimental signature of programmable quantum annealing," Nat. Commun. **4**, 2067 (2013).

[13] Tameem Albash, Walter Vinci, Anurag Mishra, Paul A. Warburton, and Daniel A. Lidar, "Consistency tests of classical and quantum models for a quantum annealer," Phys. Rev. A **91**, 042314– (2015).

[14] Mohammad H. Amin, "Searching for quantum speedup in quasistatic quantum annealers," Phys. Rev. A **92**, 052323 (2015).

[15] Steven H. Adachi and Maxwell P. Henderson, "Application of quantum annealing to training of deep neural networks," arXiv:1510.06356 (2015).

[16] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, "Quantum-assisted learning of graphical models with arbitrary pairwise connectivity," ArXiv e-prints (2016), arXiv:1609.02542 [quant-ph].

[17] "Quantum enhanced optimization (qeo)," .

[18] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, L. M. Johnson, M. A. Gouker, and W. D. Oliver, "Fabrication process and properties of fully-planarized deep-submicron nb/al-alox/nb josephson junctions for vlsi circuits," IEEE Transactions on Applied Superconductivity **25**, 1–12 (2015).

[19] S. K. Tolpygo, V. Bolkhovsky, T. J. Weir, C. J. Galbraith, L. M. Johnson, M. A. Gouker, and V. K. Semenov, "Inductance of circuit structures for mit ll superconductor electronics fabrication process with 8 niobium layers," IEEE Transactions on Applied Superconductivity **25**, 1–5 (2015).

[20] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "Thermal and residual excited-state population in a 3d transmon qubit," Phys. Rev. Lett. **114**, 240501 (2015).

[21] "D-wave systems previews 2000-qubit quantum system," .

[22] M W Johnson, P Bunyk, F Maibaum, E Tolkacheva, A J Berkley, E M Chapple, R Harris, J Johansson, T Lanting, I Perminov, E Ladizinsky, T Oh, and G Rose, "A scalable control system for a superconducting adiabatic quantum optimization processor," Superconductor Science and Technology **23**, 065004 (2010).

[23] A J Berkley, M W Johnson, P Bunyk, R Harris, J Johansson, T Lanting, E Ladizinsky, E Tolkacheva, M H S Amin, and G Rose, "A scalable readout system for a superconducting adiabatic quantum optimization system," Superconductor Science and Technology **23**, 105014 (2010).

[24] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, "Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor," Phys. Rev. B **82**, 024511 (2010).

[25] P. I Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, AJ. Berkley, R. Harris, J. P. Hilton, T. Lanting, AJ. Przybysz, and J. Whittaker, "Architectural considerations in the design of a superconducting quantum annealing processor," IEEE Transactions on Applied Superconductivity **24**, 1–10 (Aug. 2014).

[26] J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, and C. C. McGeoch, "Benchmarking a quantum annealing processor with the time-to-target metric," ArXiv e-prints (2015), arXiv:1508.05087 [quant-ph].

[27] T. Kato, "On the adiabatic theorem of quantum mechanics," J. Phys. Soc. Jap. **5**, 435 (1950).

[28] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler, "Bounds for the adiabatic approximation with applications to quantum computation," J. Math. Phys. **48**, – (2007).

[29] Lorenzo Campos Venuti, Tameem Albash, Daniel A. Lidar, and Paolo Zanardi, "Adiabaticity in open quantum systems," arXiv:1508.05558 (2015).

[30] This is equivalent to having a bounded degree connectivity graph.

[31] The analysis is based on the equivalence between the Canonical and the Microcanonical Ensembles of Statistical Mechanics. This equivalence is reviewed in many places, see e.g., Ref. [67].

[32] Equivalently, $Ns(e)$ is the logarithm of the number of

microstates with intensive energy e.
[33] A relation best known as the second law of thermodynamics $T\mathrm{d}e = \mathrm{d}s$.
[34] Vicky Choi, "Minor-embedding in adiabatic quantum computation: I. The parameter setting problem," Quant. Inf. Proc. **7**, 193–209 (2008).
[35] Vicky Choi, "Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design," Quant. Inf. Proc. **10**, 343–353 (2011).
[36] The reader is referred to the Supplemental Information (SI) for further details.
[37] Details of these instances as well as similar results obtained for other problem classes are given in the SI.
[38] R. Saket, "A PTAS for the Classical Ising Spin Glass Problem on the Chimera Graph Structure," ArXiv e-prints (2013), arXiv:1306.6943 [cs.DS].
[39] By no means however, is it meant that PTAS is able to thermally sample from a Boltzmann distribution of the input problem. In fact, it should be clear that PTAS does not.
[40] C. J. Geyer, "Parallel tempering," in *Computing Science and Statistics Proceedings of the 23rd Symposium on the Interface*, edited by E. M. Keramidas (American Statistical Association, New York, 1991) p. 156.
[41] Koji Hukushima and Koji Nemoto, "Exchange monte carlo method and application to spin glass simulations," Journal of the Physical Society of Japan **65**, 1604–1608 (1996), http://dx.doi.org/10.1143/JPSJ.65.1604.
[42] Details of our PT implementation can be found in the SI.
[43] Similar scaling behavior for other classes of Hamiltonians, specifically 3-regular 3-XORSAT instances and random ± 1 instances, is also observed, and we give the results for these instances in the SI.
[44] T. Jörg, J. Lukic, E. Marinari, and O. C. Martin, "Strong universality and algebraic scaling in two-dimensional ising spin glasses," Phys. Rev. Lett. **96**, 237205 (2006).
[45] C. K. Thomas, D. A. Huse, and A. A. Middleton, "Zero and low temperature behavior of the two-dimensional $\pm j$ ising spin glass," Phys. Rev. Lett. **107**, 047203 (2011), arXiv:1103.1946.
[46] Francesco Parisen Toldin, Andrea Pelissetto, and Ettore Vicari, "Finite-size scaling in two-dimensional ising spin-glass models," Phys. Rev. E **84**, 051116 (2011).
[47] L. A. Fernandez, E. Marinari, V. Martin-Mayor, G. Parisi, and J. J. Ruiz-Lorenzo, "Universal critical behavior of the two-dimensional ising spin glass," Phys. Rev. B **94**, 024402 (2016).
[48] A. Lucas, "Ising formulations of many NP problems," Front. Phys. **2**, 5 (2014).
[49] Walter Vinci and Daniel A. Lidar, "Optimally stopped optimization," Phys. Rev. Applied **6**, 054016 (2016).
[50] Sergio Boixo, Troels F. Ronnow, Sergei V. Isakov, Zhihui Wang, David Wecker, Daniel A. Lidar, John M. Martinis, and Matthias Troyer, "Evidence for quantum annealing with more than one hundred qubits," Nat. Phys. **10**, 218–224 (2014).
[51] Victor Martin-Mayor and Itay Hen, "Unraveling quantum annealers using classical hardness," Scientific Reports **5**, 15324 EP – (2015).
[52] Zheng Zhu, Andrew J. Ochoa, Stefan Schnabel, Firas Hamze, and Helmut G. Katzgraber, "Best-case performance of quantum annealers on native spin-glass benchmarks: How chaos can affect success probabilities," Phys. Rev. A **93**, 012317 (2016).
[53] Kristen L Pudenz, Tameem Albash, and Daniel A Lidar, "Error-corrected quantum annealing with hundreds of qubits," Nat. Commun. **5**, 3243 (2014).
[54] Kristen L. Pudenz, Tameem Albash, and Daniel A. Lidar, "Quantum annealing correction for random Ising problems," Phys. Rev. A **91**, 042302 (2015).
[55] Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A. Lidar, "Quantum annealing correction with minor embedding," Phys. Rev. A **92**, 042310–(2015).
[56] Shunji Matsuura, Hidetoshi Nishimori, Tameem Albash, and Daniel A. Lidar, "Mean field analysis of quantum annealing correction," arXiv:1510.07709 (2015).
[57] Walter Vinci, Tameem Albash, and Daniel A Lidar, "Nested quantum annealing correction," Npj Quantum Information **2**, 16017 EP – (2016).
[58] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, Tameem Albash, and Daniel A. Lidar, "Quantum-annealing correction at finite temperature: Ferromagnetic p-spin models," Phys. Rev. A **95**, 022308 (2017).
[59] S. P. Jordan, E. Farhi, and P. W. Shor, "Error-correcting codes for adiabatic quantum computation," Phys. Rev. A **74**, 052322 (2006).
[60] Adam D. Bookatz, Edward Farhi, and Leo Zhou, "Error suppression in hamiltonian-based quantum computation using energy penalties," Physical Review A **92**, 022317–(2015).
[61] Zhang Jiang and Eleanor G. Rieffel, "Non-commuting two-local hamiltonians for quantum error suppression," Quantum Information Processing **16**, 89 (2017).
[62] Milad Marvian and Daniel A. Lidar, "Error suppression for hamiltonian-based quantum computation using subsystem codes," Phys. Rev. Lett. **118**, 030504 (2017).
[63] C.H. Papadimitriou, *Computational Complexity* (Addison Wesley Longman, Reading, Massachusetts, 1995).
[64] Leslie G. Valiant, "The complexity of enumeration and reliability problems," SIAM Journal on Computing **8**, 410–421 (1979).
[65] Vilhelm Dahllöf, Peter Jonsson, and Magnus Wahlström, "Counting models for 2sat and 3sat formulae," Theoretical Computer Science **332**, 265 – 291 (2005).
[66] L. Campos Venuti, T. Albash, M. Marvian, D. Lidar, and P. Zanardi, "Relaxation vs. adiabatic quantum steady state preparation: which wins?" ArXiv e-prints (2016), arXiv:1612.07979 [quant-ph].
[67] V. Martín-Mayor, "Microcanonical approach to the simulation of first-order phase transitions," Phys. Rev. Lett. **98**, 137207 (2007).
[68] Andrew D. King, Trevor Lanting, and Richard Harris, "Performance of a quantum annealer on range-limited constraint satisfaction problems," arXiv:1502.02098 (2015).
[69] Firas Hamze and Nando de Freitas, "From fields to trees," in *UAI*, edited by David Maxwell Chickering and Joseph Y. Halpern (AUAI Press, Arlington, Virginia, 2004) pp. 243–250.
[70] Alex Selby, "Efficient subgraph-based sampling of ising-type models with frustration," arXiv:1409.3934 (2014).
[71] Lov K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," Phys. Rev. Lett. **79**, 325–328 (1997).
[72] Jérémie Roland and Nicolas J. Cerf, "Quantum search by local adiabatic evolution," Phys. Rev. A **65**, 042308–(2002).

Supplemental Information for "Temperature scaling law for quantum annealing optimizers"

THE DW2X EXPERIMENTAL QUANTUM ANNEALING OPTIMIZER

Description of the processor

The experimental results shown in the main text were taken on a 3rd generation D-Wave processor, the DW2X 'Washington' processor, installed at the Information Sciences Institute - University of Southern California (ISI). The processor connectivity is given by a 12×12 grid of unit cells, where each unit cell is composed of 8 qubits with a $K_{4,4}$ bipartite connectivity, forming the 'Chimera' graph [34, 35] with a total of 1152 qubits. Due to miscalibration, there are only 1098 operational qubits on the ISI machine. This is illustrated in Fig. 4.

The device implements the quantum annealing protocol given by the time-dependent Hamiltonian:

$$H_{\rm QA}(s) = A(s)H_{\rm D} + B(s)H \tag{10}$$

where $H_{\rm D} = -\sum_i \sigma_i^x$ is the standard transverse field driver Hamiltonian, H is the Ising Hamiltonian [Eq. (1) of the main text], and $A(s), B(s)$ are the annealing schedules satisfying $A(0) \gg B(0)$, $A(1) \ll B(1)$, and $s \equiv t/t_f \in [0,1]$ is the dimensional time annealing parameter. The predicted functional form for these schedules is shown in Fig. 5.

Details of the experiment and additional results

The randomly generated instances tested on the D-Wave processor were run with 20 random gauges [12] with 5000 reads per gauge/cycle for a total of $100,000$ anneals per instance. The annealing time chosen for the runs was the default 20μ-sec. We further corroborated the analytical derivations discussed in the main text using experiments on the commercial DW2X processor on randomly generated bi-modal $J_{ij} = \pm 1$ instances. As with the planted-solution instances, we first generate random instances of differently sized sub-graphs of the DW2X Chimera connectivity graph [34, 35] and run them multiple times on the annealer, recording the obtained energies. Figure 6 depicts the resultant residual energy $(E - E_0)$ distributions of a typical instance. As is evident, increasing the problem size N 'pushes' the energy distribution farther and farther away from the ground state value, as well as broadening the distribution and making it more gaussian-like. In the inset we measure the departure of $\langle H \rangle_\beta$ from E_0 and the spread of the energies $\sigma_\beta(H)$ over 100 random bi-modal instances per sub-graph size as a function of problem size N. For sufficiently large problem sizes, we find that the scaling of $\langle H - E_0 \rangle_\beta$ is linear while $\sigma_\beta(H)$ scales as $\sqrt{N}$, consistent with the analytical prediction (the solid lines are linear and square-root fits, respectively, to the largest four sizes).

SIMULATION METHODS

Instance generation

For the generation of instances in this work we have chosen one problem class to be that of the 'planted solution' type—an idea borrowed from constraint satisfaction (SAT) problems. In this problem class, the planted solution represents a ground-state configuration of the Hamiltonian that minimizes the energy and is known in advance. The Hamiltonian of a planted-solution spin glass is a sum of terms, each of which consists of a small number of connected spins, namely, $H = \sum_j H_j$[11]. Each term H_j is chosen such that one of its ground-states is the planted solution. It follows then that the planted solution is also a ground-state of the total Hamiltonian, and its energy is the ground-state energy of the Hamiltonian. Knowing the ground-state energy in advance circumvents the need to verify the ground-state energy using exact (provable) solvers, which rapidly become too expensive computationally as the number of variables grows. The interested reader will find a more detailed discussion of planted Ising problems in Refs.[11, 68].

For the random ± 1 instances on Chimera, we randomly (with equal probability) assign a value ± 1 to all the edges of the Chimera graph. While the ground state energy for these instances is not known with 100% certainty, we ran the Hamze-Freitas-Selby algorithm (HFS) [69, 70] for a sufficiently long time such that we were confident of having found the ground state for these instances.

For the 3-regular 3-XORSAT instances, for each spin, we randomly pick three other spins to which to couple. All couplings are picked to be antiferromagnetic with strength 1. Because all terms in the Hamiltonian are of the form $+\sigma_i^z\sigma_j^z\sigma_k^z$, the ground state is simply that all-spins-down state.

Parallel tempering

For the planted-solution instances, we first 'warmed-up' our parallel tempering simulation with 5×10^5 (for the smaller sizes) to 2×10^6 (for the larger sizes) swaps with 10 Monte Carlo sweeps per swap. The temperature distribution is picked as follows:

$$\beta_i = \left(\frac{\beta_{63}}{\beta_0}\right)^{i/63} \beta_0 \ , \quad i = 0, 1, \ldots, 63 \tag{11}$$

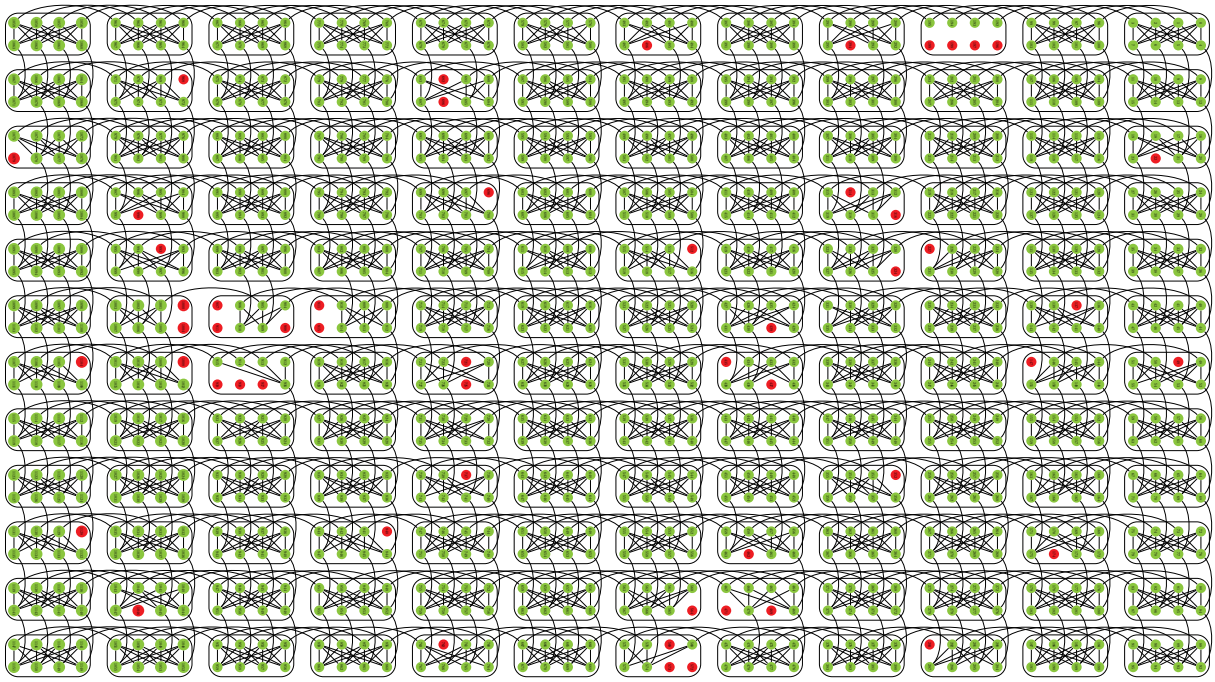

FIG. 4. **A visualization of the DW2X graph.** Operational qubits are shown in green, and inoperationable ones are shown in red. Programmable couplers are shown as black lines connecting the qubits.

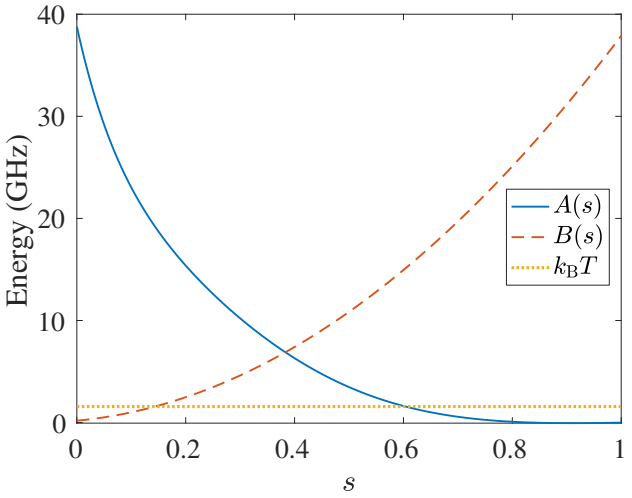

FIG. 5. **The DW2X annealing schedules.** Energy units are for $\hbar = 1$, and the operating temperature of $12m$K is shown as well.

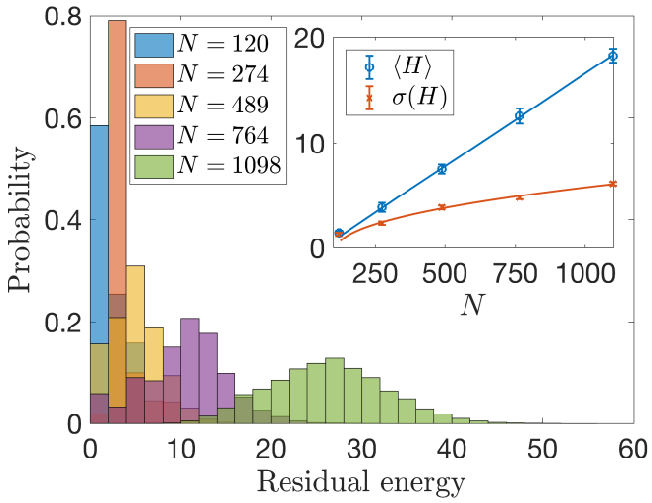

FIG. 6. **Distributions of residual energy, $E - E_0$, from DW2X simulations on random ± 1 instances.** As problem sizes grow, the distributions become more Gaussian-like. **Inset:** Gaussians' mean (blue) and standard deviation (red) as a function of problem size, averaged over 100 instances per size. The solid lines are fits to a straight line and a square root, respectively, taking into accounts all sizes but the smallest.

with $\beta_0 = 20$ and $\beta_{63} = 0.1$. After the warm-up, we sample the energy after every 50 swaps in order to minimize correlation between the energies. We use a total of 10^4 sample points, from which we extract the energies at different quantiles. In order to ensure that we have reached a thermal or near-thermal distribution, we performed the following check. The 10^4 sample points are divided into three blocks: (a) 5×10^3 samples from the last half of the samples; (b) 2.5×10^3 samples from the second quarter of the samples; (c) 1.25×10^3 samples from the second eighth of the samples. We then calculated the specific heat using the samples from each block separately; if the system has sufficiently thermalized and the samples are sufficiently uncorrelated, we expect to observe no change in the specific heat for the three sets of samples within

the error bars. We show the results of this test in Fig. 7, where we indeed observe no significant difference.

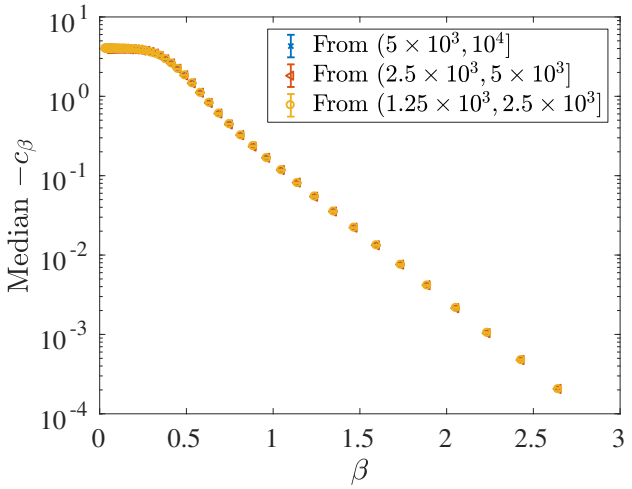

FIG. 7. **The behavior of the median specific heat for the planted-solution instances at $L = 22$ using different blocks of the samples**. Of the total of 10^4 samples, different partitions (as indicated by the legend) are used to calculate the specific heat.

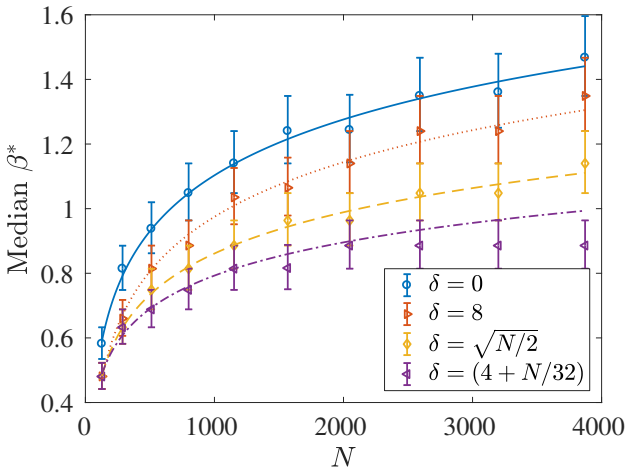

FIG. 8. **The behavior of the median value (over 100 instances) of the minimum inverse-temperature required such that for $q = 10^{-3}$ the target energy is the ground state.** Error bars correspond to the spacing between the β values of the PT simulations. Lines correspond to the fits $\beta = a + b \ln L$ with $b = 0.4990 \pm 0.1036, 0.4880 \pm 0.0922, 0.3685 \pm 0.0846, 0.2967 \pm 0.2547$ for $\delta = 0, 8, \sqrt{N/2}, (4 + N/32)$ respectively, with the uncertainty representing the 95% confidence interval for the fit parameters.

RESULTS FOR PLANTED-SOLUTION INSTANCES WITH A TARGET ENERGY

In Fig. 8, we supplement the results presented in the main text with the scaling of β when the target energy need not be the ground state, specifically $E_T = E_0 + \delta$. We consider three cases: (i) a constant about the ground state, $E_T = E_0 + 8$, (ii) a square-root scaling above the ground state, $E_T = E_0 + \sqrt{N/2}$, and a linear scaling above the ground state $E_T = E_0 + (4 + N/32)$. The specific values were picked so that the three cases would have the same target energy at the smallest size of $N = 128$. If we fit all curves with a logarithmic dependence on N, we observe a similar scaling for the cases of $\delta =$ constant, and the case of $\delta \propto \sqrt{N}$ still exhibits a logarithmic scaling but with a milder coefficient. For the case of $\delta \propto N$, the required β approaches a constant for sufficiently large problem sizes.

RESULTS FOR THE 3-REG 3XORSAT AND RANDOM ±1 CHIMERA INSTANCES

Here we provide the equivalent plots to Fig. 3 of the main text but for the 3-regular 3-XORSAT (Fig. 10) and random ±1 instances (Fig. 11). The random ±1 instances were warmed-up with up to 24×10^6 PT swaps depending on their size, while the XORSAT instances were warmed-up for with up to 200×10^6 swaps depending on their size. For both, as in the planted-solution case, 10^4 samples were taken with one sample after every 50 PT swaps. We perform the same thermalization test as for the planted-solution instances, and we observe no significant difference for the different blocks of samples (see Fig. 9).

We note that for both of these classes of instances, the β values required fall in the regime where the scaling of the specific heat with β is not yet exponential. The scaling behavior of β^* is consistent with both a $\log N$ and a $N^{1/\alpha}$ behavior.

SCALING LAWS FOR TEMPERATURES: ANALYTICAL EXAMPLES

Let us consider the simple case of non-interacting spins in a global magnetic field. This case is particularly relevant if the initial state of the quantum annealer is prepared as the thermal state of the standard driver Hamiltonian $-\sum_{i=1}^{N} \sigma_i^x$ with no overall energy scaling. The partition function is given by:

$$Z = \sum_{k=0}^{N} \binom{N}{k} e^{-\beta(k-N/2)} = [2\cosh(\beta/2)]^N \qquad (12)$$

Note that each energy spectrum has a degeneracy that grows polynomially with N. The mean energy is given by:

$$\mu/N = -\frac{1}{2}\tanh(\beta/2) \qquad (13)$$

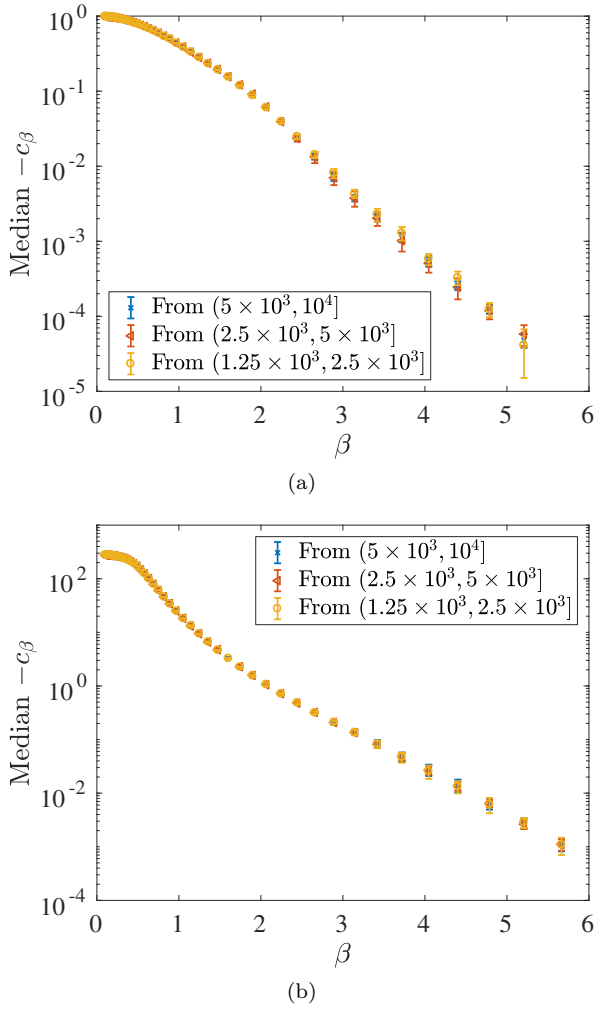

FIG. 9. **The behavior of the median specific heat for the (a) 3-regular 3-XORSAT instances ($N = 100$) and (b) the bimodal instances ($L = 12$) using different blocks of the samples**. Of the total of 10^4 samples, different partitions (as indicated by the legend) are used to calculate the specific heat.

and the standard deviation is:

$$\sigma/\sqrt{N} = \frac{1}{2}\text{sech}(\beta/2) \tag{14}$$

The ground state probability on a thermal state is then given by $p_0 = \frac{e^{\beta N/2}}{Z}$, which we can then invert to write the inverse-temperature as:

$$\beta = -\ln\left(1 - p_0^{-1/n}\right) \tag{15}$$

If we pick p_0 to be some small but fixed (independent of system size) number and take the large N limit, we find that

$$\beta = \ln(N) - \ln(-\ln p_0) + \frac{1}{2N}\ln p_0 + \dots \tag{16}$$

Therefore, we find that for this simple problem, in order to maintain a constant ground state probability while the system size grows, we must scale the inverse-temperature logarithmically with system size.

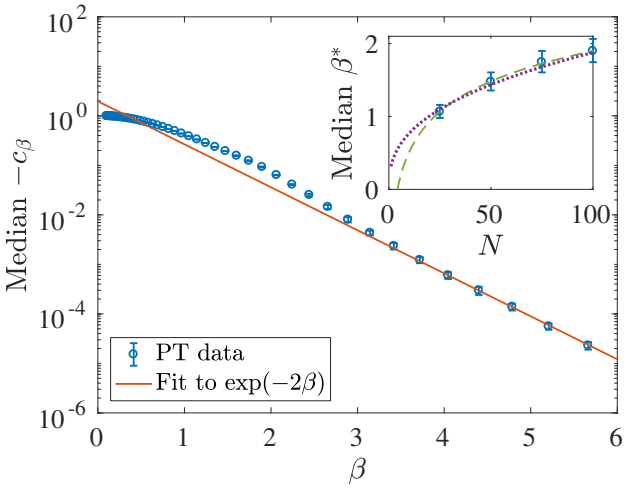

FIG. 10. **Behavior of the median specific heat (over 100 instances) for the 3-regular 3-XORSAT instances with inverse-temperature β for $N = 100$.** The behavior transitions from a polynomial scaling with β to an exponential scaling with β. **Inset:** Typical minimum inverse-temperature required for instances of size N such that probability of the ground state is at least $q = 10^{-1}$. Also shown are fits to $\log N$ anto d a power-law cN^α with $\alpha = 0.39 \pm 0.18$, which is almost indistinguishable from the logarithmic fit for large size.

A Grover search problem [71, 72] on the other hand yields the worst case scaling. In this case, we take a single state to have energy $-N$, while the remaining states have energy $-N + 1$. The partition function is given by:

$$Z = e^{\beta N}\left[1 + (2^N - 1)e^{-\beta}\right] \tag{17}$$

with mean energy:

$$\mu = -N + 1 - \frac{e^\beta}{2^N - 1 + e^\beta} \tag{18}$$

and standard deviation

$$\sigma = e^{\beta/2}\frac{\sqrt{2^N - 1}}{2^N - 1 + e^\beta} \tag{19}$$

Unlike our other local example the σ does not scale as $\sqrt{N}$. The ground state probability is given by $p_0 = 1/Z$. Inverting this for β, we find:

$$\beta = \ln(2^N - 1) - \ln\left(p_0^{-1} - 1\right) \tag{20}$$

Again, for a fixed and small p_0, expanding for large N, we get:

$$\beta = N\ln 2 - \ln\left(p_0^{-1} - 1\right) - 2^{-N} + \dots \tag{21}$$

Therefore, in this case, β must grow linearly with N in order to maintain a constant p_0. Note of course that the Grover Hamiltonian is highly non-local as it contains N-body terms.

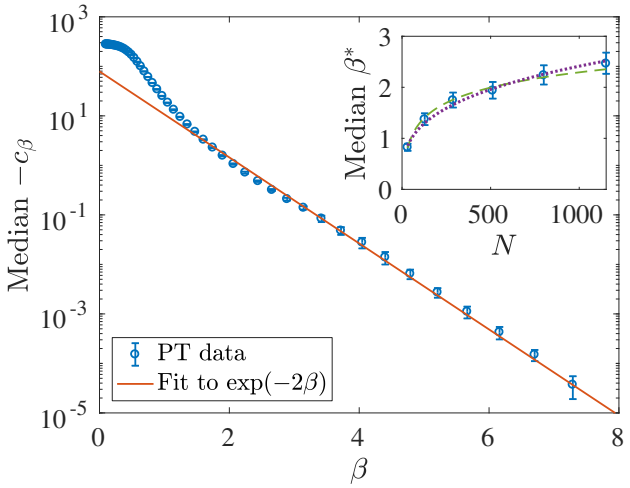

FIG. 11. **Behavior of the median specific heat (over 100 instances) for the random ± 1 instances with inverse-temperature β for $N = 1152$.** The behavior transitions from a polynomial scaling with β to an exponential scaling with β. **Inset:** Typical minimum inverse-temperature required for instances of size N such that probability of the target energy $E_{\rm T} = E_0 + \delta(N)$ is at least $q = 10^{-1}$. Also shown are fits to $\log N$ for all three cases and to a power-law cN^α with $\alpha = 0.30 \pm 0.09$, which is almost indistinguishable from the logarithmic fit for large sizes.